


[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)

 Search: ☒ The ACM Digital Library ☐ The Guide

lock manager and parent lock and count and nodes



THE ACM DIGITAL LIBRARY


[Feedback](#) [Report a problem](#) [Satisfaction survey](#)

Terms used

lock manager and parent lock and count and nodes

Found 27,465 of 185,942

Sort results by

relevance

Display results

expanded form

[Save results to a Binder](#) [Search Tips](#)☐ Open results in a new window[Try an Advanced Search](#)[Try this search in The ACM Guide](#)

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

Relevance scale ☐ ☐ ☐ ☐ ☐**1 [Query processing for XML data: Locking-aware structural join operators for XML](#)**[query processing](#)

Christian Mathis, Theo Härder, Michael Haustein

June 2006 **Proceedings of the 2006 ACM SIGMOD international conference on Management of data SIGMOD '06****Publisher:** ACM PressFull text available: [pdf\(519.20 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

As observed in many publications so far, the matching of twig pattern queries (i.e., queries that contain only the child and the descendant axis) is a core operation in XML database management systems (XDBMSs) for which the structural join and the holistic twig join algorithms were proposed. In a single-user environment, especially the latter algorithm provides a good evaluation strategy. However, when it comes to multi-user access to a single XML document, it may lead to extensive blocking situ ...

Keywords: XML documents, concurrency control, node labeling, query processing, structural joins

2 [Distributed transactions for reliable systems](#)

Alfred Z. Spector, Dean Daniels, Daniel Duchamp, Jeffrey L. Eppinger, Randy Pausch

December 1985 **ACM SIGOPS Operating Systems Review , Proceedings of the tenth ACM symposium on Operating systems principles SOSP '85**, Volume 19 Issue 5**Publisher:** ACM PressFull text available: [pdf\(1.44 MB\)](#) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)**3 [Fast detection of communication patterns in distributed executions](#)**

Thomas Kunz, Michiel F. H. Seuren

November 1997 **Proceedings of the 1997 conference of the Centre for Advanced Studies on Collaborative research****Publisher:** IBM PressFull text available: [pdf\(4.21 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Understanding distributed applications is a tedious and difficult task. Visualizations based on process-time diagrams are often used to obtain a better understanding of the

execution of the application. The visualization tool we use is Poet, an event tracer developed at the University of Waterloo. However, these diagrams are often very complex and do not provide the user with the desired overview of the application. In our experience, such tools display repeated occurrences of non-trivial commun ...

4 Concurrency and recovery in generalized search trees



Marcel Kornacker, C. Mohan, Joseph M. Hellerstein

June 1997 **ACM SIGMOD Record**, **Proceedings of the 1997 ACM SIGMOD international conference on Management of data SIGMOD '97**, Volume 26 Issue 2

Publisher: ACM Press

Full text available: [pdf\(1.59 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This paper presents general algorithms for concurrency control in tree-based access methods as well as a recovery protocol and a mechanism for ensuring repeatable read. The algorithms are developed in the context of the Generalized Search Tree (GiST) data structure, an index structure supporting an extensible set of queries and data types. Although developed in a GiST context, the algorithms are generally applicable to many tree-based access methods. The concurrency control protocol is base ...

5 The hB π -tree: a multi-attribute index supporting concurrency, recovery and node consolidation

Georgios Evangelidis, David Lomet, Betty Salzberg

February 1997 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 6 Issue 1

Publisher: Springer-Verlag New York, Inc.

Full text available: [pdf\(314.15 KB\)](#) Additional Information: [full citation](#), [abstract](#), [citations](#), [index terms](#)

We propose a new multi-attribute index. Our approach combines the hB-tree, a multi-attribute index, and the π -tree, an abstract index which offers efficient concurrency and recovery methods. We call the resulting method the hB π -tree. We describe several versions of the hB π -tree, each using a different node-splitting and index-term-posting algorithm. We also describe a new node deletion algorithm. We have implemented all the versions of the hB π -tree. Our performance results ...

Keywords: Concurrency, Multi-attribute index, Node consolidation, Recovery

6 Articles: Evaluating lock-based protocols for cooperation on XML documents



Sven Helmer, Carl-Christian Kanne, Guido Moerkotte

March 2004 **ACM SIGMOD Record**, Volume 33 Issue 1

Publisher: ACM Press

Full text available: [pdf\(104.02 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#)

We discuss four different core protocols for synchronizing access to and modifications of XML document collections. These core protocols synchronize structure traversals and modifications. They are meant to be integrated into a native XML base management System (XBMS) and are based on two phase locking. We also demonstrate the different degrees of cooperation that are possible with these protocols by various experimental results. Furthermore, we also discuss extensions of these core protocols to ...

7 The architecture of the EXODUS extensible DBMS

Michael J. Carey, David J. DeWitt, Daniel Frank, M. Muralikrishna, Goetz Graefe, Joel E. Richardson, Eugene J. Shekita

September 1986 **Proceedings on the 1986 international workshop on Object-oriented database systems**

Publisher: IEEE Computer Society Press

Full text available:  [pdf\(1.68 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

With non-traditional application areas such as engineering design, image/voice data management, scientific/statistical applications, and artificial intelligence systems all clamoring for ways to store and efficiently process larger and larger volumes of data, it is clear that traditional database technology has been pushed to its limits. It also seems clear that no single database system will be capable of simultaneously meeting the functionality and performance requirements of such a diver ...


8 The performance of current B-tree algorithms



Theodore Johnson, Dennis Sasha

March 1993 **ACM Transactions on Database Systems (TODS)**, Volume 18 Issue 1

Publisher: ACM Press

Full text available:  [pdf\(2.87 MB\)](#) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#), [review](#)

Keywords: B-trees, concurrent B-trees, concurrent data structures, performance of concurrent algorithms

9 Performance of B⁺ tree concurrency control algorithms

V. Srinivasan, Michael J. Carey

October 1993 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 2 Issue 4

Publisher: Springer-Verlag New York, Inc.

Full text available:  [pdf\(2.67 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

A number of algorithms have been proposed to access B⁺-trees concurrently, but they are not well understood. In this article, we study the performance of various B⁺-tree concurrency control algorithms using a detailed simulation model of B⁺-tree operations in a centralized DBMS. Our study covers a wide range of data contention situations and resource conditions. In addition, based on the performance of the set of B⁺-tree concurrency control algorithms, ...

Keywords: B⁺-tree structures, data contention, lock modes, performance, resource conditions, simulation models, workload parameters

10 Implementation of Argus



B. Liskov, D. Curtis, P. Johnson, R. Scheifer

November 1987 **ACM SIGOPS Operating Systems Review , Proceedings of the eleventh ACM Symposium on Operating systems principles SOSP '87**, Volume 21 Issue 5

Publisher: ACM Press

Full text available:  [pdf\(1.34 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Argus is a programming language and system developed to support the construction and execution of distributed programs. This paper describes the implementation of Argus, with particular emphasis on the way we implement atomic actions, because this is where Argus differs most from other implemented systems. The paper also discusses the performance of Argus. The cost of actions is quite reasonable, indicating that action systems like Argus are practical.

11 The file system of an integrated local network


-  Paul J. Leach, Paul H. Levine, James A. Hamilton, Bernard L. Stumpf
March 1985 **Proceedings of the 1985 ACM thirteenth annual conference on Computer Science**
Publisher: ACM Press


Full text available:  pdf(1.78 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The distributed file system component of the DOMAIN system is described. The DOMAIN system is an architecture for networks of personal workstations and servers which creates an integrated distributed computing environment. The distinctive features of the file system include: objects addressed by unique identifiers (UIDs); transparent access to objects, regardless of their location in the network; the abstraction of a single level store for accessing all objects; and the layering of a network ...

12 Office-by-example: an integrated office system and database manager


-  Kyu-Young Whang, Art Ammann, Anthony Bolmarcich, Maria Hanrahan, Guy Hochgesang, Kuan-Tsae Huang, Al Khorasani, Ravi Krishnamurthy, Gary Sockut, Paula Sweeney, Vance Waddle, Moshé Zloof
October 1987 **ACM Transactions on Information Systems (TOIS)**, Volume 5 Issue 4
Publisher: ACM Press

Full text available:  pdf(2.86 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Office-by-Example (OBE) is an integrated office information system that has been under development at IBM Research. OBE, an extension of Query-by-Example, supports various office features such as database tables, word processing, electronic mail, graphics, images, and so forth. These seemingly heterogeneous features are integrated through a language feature called example elements. Applications involving example elements are processed by the database manager, an integrated ...

13 On-line reorganization in object databases

-  Mohana K. Lakhamraju, Rajeev Rastogi, S. Seshadri, S. Sudarshan
May 2000 **ACM SIGMOD Record , Proceedings of the 2000 ACM SIGMOD international conference on Management of data SIGMOD '00**, Volume 29 Issue 2
Publisher: ACM Press

Full text available:  pdf(283.91 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Reorganization of objects in an object databases is an important component of several operations like compaction, clustering, and schema evolution. The high availability requirements (24 × 7 operation) of certain application domains requires reorganization to be performed on-line with minimal interference to concurrently executing transactions.

In this paper, we address the problem of on-line reorganization in object databases, where a set of objects have to be migrated from one ...

14 Building knowledge base management systems

- John Mylopoulos, Vinay Chaudhri, Dimitris Plexousakis, Adel Shrufi, Thodoros Topologlou
December 1996 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 5 Issue 4
Publisher: Springer-Verlag New York, Inc.

Full text available:  pdf(403.22 KB)

Additional Information: [full citation](#), [abstract](#), [citations](#), [index terms](#)

Advanced applications in fields such as CAD, software engineering, real-time process control, corporate repositories and digital libraries require the construction, efficient access and management of large, shared knowledge bases. Such knowledge bases cannot be built using existing tools such as expert system shells, because these do not scale up, nor can they be built in terms of existing database technology, because such technology does not support the rich representational structure and infer ...

Keywords: Concurrency control, Constraint enforcement, Knowledge base management systems, Rule management, Storage management

15 Guardians and Actions: Linguistic Support for Robust, Distributed Programs

 Barbara Liskov, Robert Scheifler

July 1983 **ACM Transactions on Programming Languages and Systems (TOPLAS)**,
Volume 5 Issue 3

Publisher: ACM Press


Full text available:  [pdf\(1.66 MB\)](#) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

16 Intelligent database caching through the use of page-answers and page-traces

 Nabil Kamel, Roger King

December 1992 **ACM Transactions on Database Systems (TODS)**, Volume 17 Issue 4


Publisher: ACM Press

Full text available:  [pdf\(3.08 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

In this paper a new method to improve the utilization of main memory systems is presented. The new method is based on prestoring in main memory a number of query answers, each evaluated out of a single memory page. To this end, the ideas of page-answers and page-traces are formally described and their properties analyzed. The query model used here allows for selection, projection, join, recursive queries as well as arbitrary combinations. We also show how to apply the approach under update ...

Keywords: artificial intelligence, databases, page access

17 Performance of B-tree concurrency control algorithms


 V. Srinivasan, Michael J. Carey

April 1991 **ACM SIGMOD Record , Proceedings of the 1991 ACM SIGMOD international conference on Management of data SIGMOD '91**, Volume 20 Issue 2

Publisher: ACM Press

Full text available:  [pdf\(1.09 MB\)](#) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

18 The Recovery Manager of the System R Database Manager


 Jim Gray, Paul McJones, Mike Blasgen, Bruce Lindsay, Raymond Lorie, Tom Price, Franco Putzolu, Irving Traiger

June 1981 **ACM Computing Surveys (CSUR)**, Volume 13 Issue 2

Publisher: ACM Press

Full text available:  [pdf\(1.75 MB\)](#) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

19 A framework for the performance analysis of concurrent B-tree algorithms

 Theodore Johnson, Dennis Shasha

April 1990 **Proceedings of the ninth ACM SIGACT-SIGMOD-SIGART symposium on Principles of database systems**

Publisher: ACM Press

Full text available:  [pdf\(1.46 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Many concurrent B-tree algorithms have been proposed, but they have not yet been satisfactorily analyzed. When transaction processing systems require high levels of concurrency, a restrictive serialization technique on the B-tree index can cause a bottleneck. In this paper, we present a framework for constructing analytical performance models of concurrent B-tree algorithms. The models can predict the response time and maximum throughput. We analyze three algorithms: Naive Lock-coupling, Op ...

20 System R: relational approach to database management



M. M. Astrahan, M. W. Blasgen, D. D. Chamberlin, K. P. Eswaran, J. N. Gray, P. P. Griffiths, W. F. King, R. A. Lorie, P. R. McJones, J. W. Mehl, G. R. Putzolu, I. L. Traiger, B. W. Wade, V. Watson

June 1976 **ACM Transactions on Database Systems (TODS)**, Volume 1 Issue 2

Publisher: ACM Press

Full text available: pdf(3.18 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

System R is a database management system which provides a high level relational data interface. The systems provides a high level of data independence by isolating the end user as much as possible from underlying storage structures. The system permits definition of a variety of relational views on common underlying data. Data control features are provided, including authorization, integrity assertions, triggered transactions, a logging and recovery subsystem, and facilities for maintaining ...

Keywords: authorization, data structures, database, index structures, locking, nonprocedural language, recovery, relational model

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2006 ACM, Inc.
[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads: [Adobe Acrobat](#) [QuickTime](#) [Windows Media Player](#) [Real Player](#)



USPTO

[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)

 Search: ☒ The ACM Digital Library ☐ The Guide

parent lock and child lock and nodes



THE ACM DIGITAL LIBRARY


[Feedback](#) [Report a problem](#) [Satisfaction survey](#)
Terms used parent lock and child lock and nodes

Found 22,423 of 185,942

Sort results by

relevance



Save results to a Binder

Display results

expanded form



Search Tips

☐ Open results in a new windowTry an [Advanced Search](#)Try this search in [The ACM Guide](#)

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

Relevance scale ☐ ☐ ☐ ☐ ☐**1** [Query processing for XML data: Locking-aware structural join operators for XML](#)[query processing](#)

Christian Mathis, Theo Härder, Michael Haustein

June 2006 **Proceedings of the 2006 ACM SIGMOD international conference on Management of data SIGMOD '06**

Publisher: ACM Press

Full text available: pdf(519.20 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

As observed in many publications so far, the matching of twig pattern queries (i.e., queries that contain only the child and the descendant axis) is a core operation in XML database management systems (XDBMSs) for which the structural join and the holistic twig join algorithms were proposed. In a single-user environment, especially the latter algorithm provides a good evaluation strategy. However, when it comes to multi-user access to a single XML document, it may lead to extensive blocking situ ...

Keywords: XML documents, concurrency control, node labeling, query processing, structural joins

2 [Operation specific locking in B-trees](#)

A. Biliris

June 1987 **Proceedings of the sixth ACM SIGACT-SIGMOD-SIGART symposium on Principles of database systems**

Publisher: ACM Press

Full text available: pdf(1.22 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

B-trees have been used as an access and for both primary and secondary indexing for quite some time. This paper presents a deadlock free locking mechanism in which different processes make use of different lock types in order to reach the leaf nodes. The compatibility relations among locks on a node, do not exclusively depend on their type, but also on the node status and the number and kind of processes acting currently on the node. As a result, a number of insertion or deletion processes ...

3 [The hB \\$^{\wedge}\\$Pi\\$-tree: a multi-attribute index supporting concurrency, recovery and node consolidation](#)

Georgios Evangelidis, David Lomet, Betty Salzberg

February 1997 **The VLDB Journal — The International Journal on Very Large Data**

Bases, Volume 6 Issue 1

Publisher: Springer-Verlag New York, Inc.

Full text available:  [pdf\(314.15 KB\)](#) Additional Information: [full citation](#), [abstract](#), [citations](#), [index terms](#)

We propose a new multi-attribute index. Our approach combines the hB-tree, a multi-attribute index, and the Π -tree, an abstract index which offers efficient concurrency and recovery methods. We call the resulting method the hB Π -tree. We describe several versions of the hB Π -tree, each using a different node-splitting and index-term-posting algorithm. We also describe a new node deletion algorithm. We have implemented all the versions of the hB Π -tree. Our performance results ...

Keywords: Concurrency, Multi-attribute index, Node consolidation, Recovery

4 Articles: Evaluating lock-based protocols for cooperation on XML documents



Sven Helmer, Carl-Christian Kanne, Guido Moerkotte

March 2004 **ACM SIGMOD Record**, Volume 33 Issue 1

Publisher: ACM Press

Full text available:  [pdf\(104.02 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#)

We discuss four different core protocols for synchronizing access to and modifications of XML document collections. These core protocols synchronize structure traversals and modifications. They are meant to be integrated into a native XML base management System (XBMS) and are based on two phase locking. We also demonstrate the different degrees of cooperation that are possible with these protocols by various experimental results. Furthermore, we also discuss extensions of these core protocols to ...

5 A locking protocol for resource coordination in distributed databases



Daniel A. Menasce, Gerald J. Popek, Richard R. Muntz

June 1980 **ACM Transactions on Database Systems (TODS)**, Volume 5 Issue 2

Publisher: ACM Press

Full text available:  [pdf\(2.69 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

A locking protocol to coordinate access to a distributed database and to maintain system consistency throughout normal and abnormal conditions is presented. The proposed protocol is robust in the face of crashes of any participating site, as well as communication failures. Recovery from any number of failures during normal operation or any of the recovery stages is supported. Recovery is done in such a way that maximum forward progress is achieved by the recovery procedures. Integration of ...

Keywords: concurrency, consistency, crash recovery, distributed databases, locking protocol

6 Locking objects and classes in multiversion object-oriented databases



Wojciech Cellary, Waldemar Wiczerzycki

December 1993 **Proceedings of the second international conference on Information and knowledge management**

Publisher: ACM Press

Full text available:  [pdf\(1.09 MB\)](#) Additional Information: [full citation](#), [references](#), [index terms](#)

7 Investigations in tree locking for compiled database applications

Heng Yu, Grant E. Weddell

October 2004 **Proceedings of the 2004 conference of the Centre for Advanced Studies**

on Collaborative research**Publisher:** IBM PressFull text available:  pdf(273.86 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

We report on initial research in tree locking (TL) schemes for compiled database applications. Such applications have a repository style of architecture in which a collection of software modules operate on a common database in terms of a set of predefined transaction types, an architectural view that is also useful for embedded control programs. Since TL schemes are deadlock free, it becomes possible to entirely decouple concurrency control from any functionality relating to recovery. This pr ...

8 Edge locks and deadlock avoidance in distributed systems


Henry F. Korth

August 1982 **Proceedings of the first ACM SIGACT-SIGOPS symposium on Principles of distributed computing****Publisher:** ACM PressFull text available:  pdf(923.82 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Two locking protocols are defined for distributed database systems. One protocol provides deadlock avoidance without the need to roll back transactions. The other allows a useful weakening of the protocol in which only a limited class of easily handled deadlocks may occur. The protocols are capable of handling replicated as well as partitioned data. Like the centralized protocol on which they are based, the protocol of this paper permits locking at multiple granularities and allows collecti ...

9 Semantic locking in object-oriented database systems

Rodolfo F. Resende, Divyakant Agrawal, Amr El Abbadi

October 1994 **ACM SIGPLAN Notices , Proceedings of the ninth annual conference on Object-oriented programming systems, language, and applications OOPSLA '94**, Volume 29 Issue 10**Publisher:** ACM PressFull text available:  pdf(1.66 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Object-oriented databases are being increasingly used to model non-standard applications that emphasize modularity, composition, and rapid prototyping. A semantic locking protocol is presented for transaction management for such object-oriented databases. In particular, the protocol incorporates the semantics of complex objects, nested executions and dynamic conflicts resulting from referentially shared objects.

10 Concurrency and recovery for index trees

David Lomet, Betty Salzberg

August 1997 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 6 Issue 3**Publisher:** Springer-Verlag New York, Inc.Full text available:  pdf(168.36 KB) Additional Information: [full citation](#), [abstract](#), [citations](#), [index terms](#)

Although many suggestions have been made for concurrency in B⁺-trees, few of these have considered recovery as well. We describe an approach which provides high concurrency while preserving well-formed trees across system crashes. Our approach works for a class of index trees that is a generalization of the B⁺-tree. This class includes some multi-attribute indexes and temporal indexes. Structural changes in an index tree are decomposed into a sequence of atomic actions, each one ...

Keywords: Access methods, B-trees, Concurrency, Indexing, Recovery

11 Deadlock freedom using edge locks



Henry F. Korth

December 1982 **ACM Transactions on Database Systems (TODS)**, Volume 7 Issue 4

Publisher: ACM Press

Full text available: pdf(370.72 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

We define a series of locking protocols for database systems that all have three main features: freedom from deadlock, multiple granularity, and support for general collections of locking primitives. A rooted directed acyclic graph is used to represent multiple granularities, as in System R. Deadlock freedom is guaranteed by extending the System R protocol to require locks on edges of the graph in addition to the locks required on nodes.

Keywords: concurrency control, locking, serializability

12 Concurrency control and recovery for balanced B-link trees

Ibrahim Jaluta, Seppo Sippu, Eljas Soisalon-Soininen

April 2005 **The VLDB Journal – The International Journal on Very Large Data Bases**, Volume 14 Issue 2

Publisher: Springer-Verlag New York, Inc.

Full text available: pdf(302.02 KB) Additional Information: [full citation](#), [abstract](#)

In this paper we present new concurrent and recoverable B-link-tree algorithms. Unlike previous algorithms, ours maintain the balance of the B-link tree at all times, so that a logarithmic time bound for a search or an update operation is guaranteed under arbitrary sequences of record insertions and deletions. A database transaction can contain any number of operations of the form "fetch the first (or next) matching record", "insert a record", or "delete a reco ...

Keywords: Concurrency control, Recovery, Transaction, Tree-structure modifications

13 Access method concurrency with recovery



David Lomet, Betty Salzberg

June 1992 **ACM SIGMOD Record , Proceedings of the 1992 ACM SIGMOD international conference on Management of data SIGMOD '92**, Volume 21 Issue 2

Publisher: ACM Press

Full text available: pdf(981.51 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Providing high concurrency in B+-trees has been studied extensively. But few efforts have been documented for combining concurrency methods with a recovery scheme that preserves well-formed trees across system crashes. We describe an approach for this that works for a class of index trees that is a generalization of the Blink-tree. A major feature of our method is that it works with a range of different recovery methods. It achieves this by decompos ...

14 A nested transaction model for multilevel secure database management systems



Elisa Bertino, Barbara Catania, Elena Ferrari

November 2001 **ACM Transactions on Information and System Security (TISSEC)**, Volume 4 Issue 4

Publisher: ACM Press

Full text available: pdf(560.96 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

This article presents an approach to concurrency control for transactions in a Multilevel Secure Database Management System (MLS/DBMS). The major problem is that concurrency control mechanisms used in traditional DBMSs are not adequate in a

MLS/DBMS, since they may be exploited to establish covert channels. The approach presented in this article, which uses single-version data items, is based on the use of nested transactions, application-level recovery, and notification-based locking protocols.

...

Keywords: Nested transactions, concurrency control, covert channels, multilevel secure database management systems

15 Concurrency control in advanced database applications



Naser S. Barghouti, Gail E. Kaiser

September 1991 **ACM Computing Surveys (CSUR)**, Volume 23 Issue 3

Publisher: ACM Press

Full text available: [pdf\(4.69 MB\)](#) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

Keywords: advanced database applications, concurrency control, cooperative transactions, design environments, extended transaction models, long transactions, object-oriented databases, relaxing serializability

16 Uncoupling updating and rebalancing in chromatic binary search trees



Otto Nurmi, Eljas Soisalon-Soininen

April 1991 **Proceedings of the tenth ACM SIGACT-SIGMOD-SIGART symposium on Principles of database systems**

Publisher: ACM Press

Full text available: [pdf\(724.54 KB\)](#) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

17 Anatomy of a native XML base management system

T. Fiebig, S. Helmer, C.-C. Kanne, G. Moerkotte, J. Neumann, R. Schiele, T. Westmann
December 2002 **The VLDB Journal – The International Journal on Very Large Data Bases**, Volume 11 Issue 4

Publisher: Springer-Verlag New York, Inc.

Full text available: [pdf\(300.97 KB\)](#) Additional Information: [full citation](#), [abstract](#), [citations](#), [index terms](#)

Several alternatives to manage large XML document collections exist, ranging from file systems over relational or other database systems to specifically tailored XML base management systems. In this paper we give a tour of Natix, a database management system designed from scratch for storing and processing XML data. Contrary to the common belief that management of XML data is just another application for traditional databases like relational systems, we illustrate how almost every component in a ...

Keywords: Database, XML

18 Concurrency and recovery in generalized search trees



Marcel Kornacker, C. Mohan, Joseph M. Hellerstein

June 1997 **ACM SIGMOD Record , Proceedings of the 1997 ACM SIGMOD international conference on Management of data SIGMOD '97**, Volume 26 Issue 2

Publisher: ACM Press

Full text available: [pdf\(1.59 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This paper presents general algorithms for concurrency control in tree-based access

methods as well as a recovery protocol and a mechanism for ensuring repeatable read. The algorithms are developed in the context of the Generalized Search Tree (GiST) data structure, an index structure supporting an extensible set of queries and data types. Although developed in a GiST context, the algorithms are generally applicable to many tree-based access methods. The concurrency control protocol is base ...

19 A symmetric concurrent B-tree algorithm

Vladimir Lanin, Dennis Shasha

November 1986 **Proceedings of 1986 ACM Fall joint computer conference**

Publisher: IEEE Computer Society Press

Full text available:  pdf(1.10 MB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

20 Nested transactions and read-write locking



A. Fekete, N. Lynch, M. Merrit, W. Weihl

June 1987 **Proceedings of the sixth ACM SIGACT-SIGMOD-SIGART symposium on Principles of database systems**

Publisher: ACM Press

Full text available:  pdf(1.93 MB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2006 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)

Patent Storm

[Home](#) [Browse by Inventor](#) [Browse by Date](#) [Resources](#) [Contact Us](#)

Type your search term here



[Patents by Class](#)

United States Patent Class 710/200 ACCESS LOCKING

[Class Hierarchy](#)

Patent Infringement - No upfront fees makes the damage recovery decision easy www.Patent-Infringement.org

[Ads by Google](#)

[Adv](#)

[Previous](#) |

Patent No. **Patent Title:**

[7114157](#)

System controlling exclusive access by control programs to system resources

The present invention is a system for controlling exclusive access by control programs (30) to system resources (52). A system in accordance with the invention includes a plurality of addressable locations (12, 14, 16, 18, 20 and 22) in the system; a communication system (24) connecting the...

[7114013](#)

Intelligent data storage manager

The intelligent data storage manager functions to combine the non-homogeneous physical devices contained in a data storage subsystem to create a logical device with new and unique quality of service characteristics that satisfy the criteria for the policies appropriate for the present data...

[7107377](#)

Information processing apparatus and method of controlling an access level

An information processing apparatus includes a process part implementing an access to a designated address having an access level, an access detector detecting the access, and an access part being capable of changing the access level of the designated address when the access is detected by the...

[7106169](#)

Methods and apparatus for verifying the installation of components in a system

A verification system for an interchangeable component configured to be mated with a receiving system. The system includes a key device supported by the interchangeable component. The key device includes a transmitter configured to transmit a signal comprising component identification...

[7103887](#)

Load-balancing queues employing LIFO/FIFO work stealing

In response to source code that represents instructions for dynamically allocating memory to objects, a compiler/interpreter produces instructions that implement a garbage collector. The garbage collector operates in garbage-collection cycles, which include parallel-execution operations such as...

[7103766](#)

System and method for making BIOS routine calls from different hardware...

The specification may disclose systems and related methods for ensuring that as between two partitions in a computer system, each partition using a separate operating system, calls to basic input output system (BIOS) routines are managed in such a way that only one instance of a BIOS routine is...

- 7096324 Embedded processor with dual-port SRAM for programmable logic**
Methods and apparatus for programmable logic devices including embedded processors having a dual-port SRAMs. A programmable logic integrated circuit includes a programmable logic portion having a plurality of logic elements, programmably configurable to implement user-defined combinatorial or...
- 7093230 Lock management thread pools for distributed data systems**
A distributed data system may include a plurality of nodes one or more of which may include at least one multi-threaded process operable to access portions of distributed data. A lock mechanism may grant locks to the multi-threaded processes for portions of the distributed data. Only a process...
- 7089555 Ordered semaphore management subsystem**
An ordered semaphore management subsystem and method for use in an application system which includes a plurality of processors competing for shared resources each of which is controlled by a unique semaphore. The subsystem generates an ordered semaphore field (OSF) corresponding to each...
- 7089540 Compiling method with exception handling**
When an execution program is to be compiled, based on information for the inlining of a function in this program, data concerning a lock set is generated for a lock for a resource due to the inlined function. Further, based on the lock data, the data concerning a lock count set upon the...
- 7089340 Hardware management of java threads utilizing a thread processor to manage a...**
A system for managing threads to handle transaction requests connected to input/output (I/O) subsystems to enable notification to threads to complete operations.
- 7089339 Sharing of functions between an embedded controller and a host processor**
An improved system is described for allowing an embedded controller and a host processor to share access to modules in a computer system. The shared access system of the present invention enables exclusive, one-at-a-time access by a processor to a module and concurrent access by more than one...
- 7085867 Methods and structure for SCSI2 to SCSI3 reservation protocol mapping**
Methods and associated structures for transparently mapping SCSI2 reservation protocol exchanges into corresponding SCSI3 reservation protocol exchanges. A mapping element may be operable within host systems that support SCSI2 reservation protocols. The mapping element intercepts and translates...
- 7082607 Reduced-overhead context-saving in static priority scheduled operating systems**
There are disclosed methods, computing devices and software products that seek to reduce memory/processor overheads in static priority scheduled operating systems by reducing the number of resource contexts that need to be saved and restored when running a plurality of tasks. This may be...
- 7082606 Dedicated heterogeneous node scheduling including backfill scheduling**
A method and system for job backfill scheduling dedicated heterogeneous nodes in a multi-node computing environment. Heterogeneous nodes are grouped into homogeneous node sub-pools. For each sub-pool, a free node schedule (FNS) is created so that the number of to chart the free nodes over time....
- 7080213 System and method for reducing shared memory write overhead in multiprocessor...**
A system and method for reducing shared memory write overhead in multiprocessor system. In one embodiment, a multiprocessing system implements a method comprising storing an indication of obtained store permission corresponding to a particular address in a store buffer. The indication may be,...
- 7073004 Method and data processing system for microprocessor communication in a...**
The address tenure for PCR synchronization operations is redefined to support inclusion of the synchronization data within the address tenure. The bits of a particular field within the address tenure (e.g., the address field) are re-allocated to synchronization data, which is known to be small...
- 7069366 System and method for handling resource transaction requests**
A system and method handle resource transaction requests. The system marks a waiting transaction request, having an access mode which is incompatible with the access mode of the executing transaction, with an ignore bit, so that a transaction request with an access mode compatible with the...
- 7069365 System and method for controlling multiple devices via general purpose...**

A computer system including a storage device, a first and second device driver, and a nexus driver. The storage device has multiple bit positions, and is coupled to a set of electrical terminals. A first portion of the electrical terminals is adapted for coupling to a first device, and a second...

7065761 Nonvolatile logical partition system data management

A logical partition (LPAR) computer system for managing partition configuration data is disclosed, which includes a nonvolatile memory, and a plurality of logical partitions, each running independently from the other logical partitions. The system also includes a console coupled to the computer...

7062615 Multi-channel memory access arbitration method and system

A method and system for allowing flexible control of access to a shared memory by multiple requesters. In a preferred embodiment, the invention arbitrates access to flash memory on a HBA between multiple host channels and HBA microprocessors, and eliminates contention possibilities for the flash...

7062583 Hardware semaphore intended for a multi-processor system

The invention relates to a method in a hardware semaphore lock (L1-LN) intended for a multi-processor system, which semaphore lock (L1-LN) protects a shared resource (R1-RN) in connection with the system in such a way that only a process which has reserved the semaphore lock (L1-LN) and...

7058629 System and method for detecting termination of an application instance using...

A system and method for detecting termination of an application, such as a database instance, that uses locks is described. A holding child process is started from a parent process. The holding child process connects to a monitored application instance, acquires an exclusive lock on the...

7055151 Systems and methods for multi-tasking, resource sharing and execution of...

In a multi-tasking pipelined processor, consecutive instructions are executed by different tasks, eliminating the need to purge an instruction execution pipeline of subsequent instructions when a previous instruction cannot be completed. The tasks do not share registers which store task-specific...

7051136 Entry locking for large data structures

The invention is based on the idea that for a large data structure with N entries, memory space for the locks corresponding to the entries can be saved by performing a hashing function on a value that represents an entry into a hashed value 1 to M. This hashed value is used to index the table of...

7047337 Concurrent access of shared resources utilizing tracking of request reception...

An apparatus, program product and method to manage access to a shared resource by a plurality of processes in a multithreaded computer via a collection of atomic operations that track both the order in which requests that use a shared resource are received, and the order in which processing of...

7047336 Method for blocking request to bus

A method for blocking a request to a front side bus interconnected between a central processing unit (CPU) and a control chip includes the following steps. First, a bus ownership of the control chip is assigned via a bus priority signal line. Any request from the CPU to the front side bus is...

7043580 Cluster lock server: ability to support multiple standard and proprietary...

An apparatus for and method of implementing a cluster lock processing system having a relatively large number of commodity cluster instruction processors which are managed by a highly scalable, off the shelf communication processor. Because the commodity processors have virtually no system...

7035981 Asynchronous input/output cache having reduced latency

The present invention is generally directed to a device including an asynchronous input/output (I/O) data cache. The device includes a single data storage area that is disposed in communication with both a system data bus and a I/O data bus. Similarly, the device includes an address storage area...

7032047 Method of regulating usage and/or concession eligibility via distributed list...

A method of regulating usage and/or concession eligibility in a smart card system is described herein. A card acceptance location (110) detects a presence of a smart card and determines its identification code. The card acceptance location (110) checks the identification code against a list...

7028119 Automated tool for detection of potential race condition

A method for detecting a potential race condition is provided comprising the steps of identifying a memory access in a source code segment, the source code segment being executable in any one of a plurality of execution threads; searching backward from the memory access through each of the...

7020725 Method of reserving isochronous resources in a network comprising a wireless...

A method for reserving an isochronous resource over a link between a first bus and a second bus, the link including a first interface device connected to the first bus and a second interface device connected to the second bus. The reservation of the resources over the link is transparent to the...

7017031 Method, apparatus and system for managing released promotion bits

A data processing system includes a global promotion facility containing a plurality of promotion bit fields, an interconnect, and a plurality of processing units coupled to the global promotion facility and to the interconnect. A first processing unit includes an instruction sequencing unit, an...

7013356 Methods and structure for preserving lock signals on multiple buses coupled...

Structure and methods for preserving lock requests by master devices on multiple buses each coupled to a port of a multiported device. The invention provides for arbitration among multiple ports of a multiported device to preserve the intent of a lock request to retain exclusive control of the...

7000048 Apparatus and method for parallel processing of network data on a single...

A method and apparatus for handling multiple processing streams in parallel on a single thread of a processing device. In one embodiment, a parallel processing agent includes a scheduler that multiplexes a number of processing streams, or pipelines, on a single thread of execution.

7000047 Mechanism for effectively handling livelocks in a simultaneous multithreading...

A method and multithreaded processor for handling livelocks in a simultaneous multithreaded processor. A number of instructions for a thread in a queue may be counted. A counter in the queue may be incremented if the number of instructions for the thread in the queue in a previous clock cycle is...

7000046 Standard channel I/O processor (SCIOP)

An apparatus for and method of implementing a cluster lock processing system having a relatively large number of commodity cluster instruction processors which are managed by a highly scalable, off the shelf communication processor. Because the commodity processors have virtually no system...

6996721 Flash device security method utilizing a check register

A security method for preventing accidental or unauthorized writes to a flash memory. According to one embodiment of the present invention, a BIOS program stored in a flash memory array generates a random access code when executed by a processor. A check register stores the random access code...

6985984 Multiprocessing systems employing hierarchical back-off locks

A multiprocessing system including multiple processing nodes employs various implementations of hierarchical back-off locks. A thread attempting to obtain a software lock may determine whether the lock is currently owned by a different node than the node in which the thread is...

6985983 Translating device adapter having a common command set for interfacing multiple

A translating host bus interface adapter is capable of connecting a computer system as a compute node to a storage area network. The adapter has a processor and a memory system containing firmware for execution in the processor. The host bus adapter is capable of recognizing first...

6985783 Data processing device with an indexed immediate addressing mode

A data processing device is provided with an indexed-immediate addressing mode for processing streams of data. An instruction register 900 receives an instruction for execution. Decoding circuitry 913 selects a register specified by a field in an instruction to provide an index value...

6983461 Method and system for deadlock detection and avoidance

The invention relates to a method of operating a computer system supporting multiple processes, and the intention is to avoid deadlock (an example of which is where process A owns resource A, and is waiting on resource B, but resource B is currently owned by process B, which is...

6981108 Method for locking shared resources connected by a PCI bus

A computer system according to the present invention comprises at least two CPUs; at least one shared system resource accessible to each of the CPUs; and shared memory shared by the CPUs. The shared memory has therein a resource locking table, comprising memory elements. Each of the...

6978330

Shared resource access via declarations that contain a sequence number of a...

Logic (also called "re-ordering semaphore") issues semaphore grants to access a shared resource in an order different from the order in which semaphore requests for accessing the shared resource are received. The re-ordering semaphore needs to receive a semaphore release between any...

6976107

Adaptive spin latches

An adaptive spin latch system is provided for use in a multiprocessor computer system. The spin latch system includes a run queue, a spin latch module and a wait queue. The run queue is adapted to store agent index data correlated to at least one agent in run mode. The spin latch...

6973652

Sequencing of tasks within customer service processing systems

Methods and apparatus for executing a series of sequential tasks existing within a system-level domain of a customer computing system providing provisioning of meta-level domain products are disclosed. For each product or product action (202), one or more sequence of tasks (250-272)...

6973521

Lock controller supporting blocking and non-blocking requests

A lock controller supports both blocking and non-blocking lock requests issued by processors of a processing engine when attempting to access a shared resource of an intermediate network device. The non-blocking lock controller capability provides a processor with the flexibility to...

6971102

Computer system, memory management method, storage medium and program...

The present invention provides a technique for skipping a locking process for an object in memory when a thread accesses an object that only it will access in order to reduce the load imposed on a system and to improve the overall system performance. A program executing in a computer...

6970963

Apparatus and method for synchronizing multiple accesses to common resources

A method and apparatus for providing a synchronization mechanism to control accesses to a non-cached shared resource by devices connected by a high speed interconnect is herein described. A semaphore is used to control access to a shared resource which is implemented by a pair of...

6965961

Queue-based spin lock with timeout

A queue-based spin lock with timeout allows a thread to obtain contention-free mutual exclusion in fair, FIFO order, or to abandon its attempt and time out. A thread may handshake with other threads to reclaim its queue node immediately (in the absence of preemption), or mark its...

6965893

Techniques for granting shared locks more efficiently

Techniques are disclosed for managing resources that are accessible to a plurality of entities. In one embodiment, shared locks on a resource are granted more efficiently by maintaining data that is local to (e.g. on the same node as) each entity to indicate whether an exclusive lock...

6963872

Adaptive lock escalation based on the concept of unescalatable locks

In this invention, we propose an adaptive lock escalation scheme that can significantly enhance the performance of the database management system under excessive lock requests. In existing lock escalation methods, under excessive lock requests, the system's performance degrades...

6961828

Cluster system, memory access control method, and recording medium

Description of a parallel program is facilitated in a cluster system which loosely couples a plurality of server computers each having a cluster file system. In the cluster system, an application program is executed to generate a process. Using the cluster file system, the process...

6959337

Networked system for assuring synchronous access to critical facilities

Maintenance of a network resource is regulated by providing a lock data area that indicates an ownership status of the network resource, and a lock server process to maintain the lock data area. A client that is to perform maintenance on the network resource sends a command to the...

6959351

Handling of a multi-access instruction in a data processing apparatus

The present invention provides a data processing apparatus and method for handling a multi-access instruction of the type which specifies that an access request of a first type and an access request of a second type should be performed without any intervening accesses taking place....

6952736

Object-based locking mechanism

An object-based locking mechanism is disclosed in which the requesting and granting of access to data is decoupled by the use of methods to regulate the data access. In one embodiment, a client object calls a request lock method of a server object when it desires access to data...

6938139

Method and system for data element change across multiple instances of data...

A method and system for updating and maintaining cache coherency across nodes in a cluster. The method uses a combination of read and write locks on the instances of the cache, with some embodiments using a master locking database. Periodic validations of timestamps are used to...

6938115

Method and computer device with different criticality

A computer system and method are provided for executing multiple software of different mission-criticality. The computer system includes at least one access control circuit to prevent access to memory assigned to software of greater mission-criticality by software of lesser...

6938114

Method and apparatus for managing access to a service processor

A method, apparatus, and computer implemented instructions for providing a plurality of terminals access to a service processor located within the data processing system. A first indication is returned if the service processor is unlocked in response to receiving a request from a...

6934806

Method and system for improving input/output performance by proactively flushing

A method (and system) of improving performance of a multiprocessor system, includes proactively flushing and locking an arbitrarily-sized region of memory out of caches of the multiprocessor system.

6925515

Producer/consumer locking system for efficient replication of file data

In a distributed file system the distributed storage management is made useful to a variety of applications. Multiple quality of service options are provided through locking. Three locking systems are provided. The system offers a locking system designed for sequential consistency...

6925551

Method, apparatus and system for accessing a global promotion facility through

A multiprocessor data processing system includes first and second processors coupled to an interconnect and to a global promotion facility containing at least one promotion bit field. The first processor initiates execution of a branch-type instruction to request acquisition of a...

6922745

Method and apparatus for handling locks

A method and device for determining an attribute associated with a locked load instruction and selecting a lock protocol based upon the attribute of the locked load instruction. Also disclosed is a method for concurrently executing the respective lock sequences associated with...

6922744

Communal lock processing system for multiprocessor computer system

In order to implement alternative pathways and procedures for handling a separate set of software locks, an arrangement of circuits is described. These circuits allow for generating and handling specific requests for communal software locks without additional software development...

6920514

Method, apparatus and system that cache promotion information within a processor

A data processing system includes a global promotion facility and a plurality of processing units coupled by an interconnect. At least one processing unit among the plurality of processing units includes one or more second caches having cache arrays in which instructions and...

6915392

Optimizing memory usage by vtable cloning

An arrangement is provided for optimizing memory usage through vtable cloning. When a request to acquire a shared object is received, it is first examined to see whether the shared object is currently locked via a vtable pointer stored in the shared object. If the shared object is not...

6910127

System and method for secure network provisioning by locking to prevent loading

A system and method have been provided for securely provisioning configuration data in a network-connected

integrated circuit device. The method comprises: receiving configuration data addressed to device registers; loading the configuration data in configuration registers; and,...

6907484

Method and apparatus for atomically changing selected bits within a register

One embodiment of the present invention provides a system that facilitates atomically updating selected bits within a register in a computing system. During operation, the system receives a command to update selected bits within the register. This command includes a data word and a...

6898650

Queueing method supporting multiple client accesses simultaneously

A queue includes a plurality of containers. Each container includes a lock. Clients, possibly in a multiprocessor system, can simultaneously access the queue, each client locking only a single container as needed. A manager can lock the entire queue to perform maintenance.

6892257

Exclusive access control to a processing resource

A data processing system is provided with multiple processors that share a main memory. Semaphore values associated with data elements within the memory system, including the main memory, are used to establish exclusive access permissions to those data elements. An exclusive access clear...

[Home](#) | [Browse by Inventor](#) | [Browse by Date](#) | [Resources](#) | [Contact Us](#)

© 2004-6 PatentStorm LLC. All rights reserved.